

Bandwidth-Limited Quantum Control

A finite-rate phase-reference test in the Penrose-overlap regime

<https://ignorantobserver.xyz>

Aernoud Dekker

May 2026

Version 1.7

DOI: 10.17605/OSF.IO/G5WRH

Abstract

We propose a prospective interferometric test of whether the physical reference frame that defines a measurement basis can become bandwidth-limited. The object of study is not a modification of the Schrödinger equation, a derivation of the Born rule, or a solution to Bell's theorem. It is a narrower control question: if the phase or basis reference $\theta(t)$ is maintained by a finite-rate tracking loop, does observer-relative online visibility breakdown scale with the deficit between the reference dynamics and the tracking channel? In a minimal phase-reference model, an unstable reference mode with entropy rate h_{KS} tracked through an effective channel of rate C_{eff} predicts

$$\frac{d}{dt} \ln \sigma_{\theta}^2 \approx 2(h_{\text{KS}} - C_{\text{eff}} \ln 2),$$

and hence $V_{\text{IOF}}(t) = \exp[-\frac{1}{2}\sigma_0^2 e^{2\kappa t}]$, with $\kappa = h_{\text{KS}} - C_{\text{eff}} \ln 2$. The primary experimental observable is therefore not an absolute visibility level but a breakdown time

$$t_{\text{break}} \propto \frac{1}{h_{\text{KS}} - C_{\text{eff}} \ln 2}$$

under deliberate throttling of the effective basis-tracking channel. In a standalone reference-channel demonstrator this is an operational visibility test, not a claim that an external computer's ignorance causes irreversible physical decoherence. The mesoscopic regime is the primary scientific target because the same finite-rate timescale can overlap with Penrose objective-reduction timescales while depending on different knobs: tracking bandwidth rather than mass geometry. We specify the assumptions, operational estimators, and falsification criteria.

Terminology convention. In this project, *The Ignorant Observer Framework* (IOF) refers to the full framework: finite self-knowledge, structural ignorance, basis tracking, philosophical inspiration, and the broader interpretation stack. *Bandwidth-Limited Quantum Control* (BLQC) is the technical module that operationalizes IOF's keystone claim as a finite-rate basis-tracking experiment. The claim originates in IOF; this paper states the controlled experimental discriminator.

Keystone claim. The central claim of this paper is not merely that finite controllers introduce noise. It is that, when the physical variable defining the measurement basis is itself tracked

through a finite-capacity channel, the unresolved online basis uncertainty should scale with the deficit rate

$$\kappa = h_{\text{KS}} - C_{\text{eff}} \ln 2.$$

If this scaling is confirmed only in a standalone online-control benchmark, then observer-side basis tracking becomes a calibrated variable rather than a purely external idealization. The stronger physics question is whether the same capacity-ratio coordinate remains predictive in the Penrose-overlap regime. If it is not confirmed there, the framework loses its scientific force.

Scope of the proposal. The proposed test can be stated without adopting a hidden-variable model, a metaphysical interpretation of quantum mechanics, or any claim about consciousness. A reference-channel demonstrator may reproduce and calibrate the expected finite-rate online tracking boundary, but that is not the main physics claim. The deeper BLQC test asks whether a controlled mesoscopic experiment contains an additional visibility channel whose timescale is fixed by finite-rate tracking of the basis reference in a regime not exhausted by ordinary latency, pulse-error, and offline-recoverable tracking artifacts. Broader interpretive consequences should be considered only after this operational claim is tested.

1 Introduction

This paper proposes a narrow experimental question. In an interferometer, spin experiment, or qubit platform, the measurement basis is not only a symbol in a Hilbert-space calculation. It is also a physical reference frame: a phase reference, field direction, pulse phase, local oscillator, or calibrated control axis. The experimenter normally treats this reference as externally known. The present proposal asks what happens when the physical reference itself is treated as an object that must be tracked by a finite-rate controller.

The motivation for treating the reference this way is structural. Any laboratory implementation includes an embedded controller that cannot stand outside the dynamics producing its own measurement question. The basis used to define the measurement is normally absorbed into a textbook idealization—“we know the setting”—but that idealization is itself an engineering claim about a finite tracking loop. The present paper asks what happens to quantum visibility when that idealization is deliberately stressed.

The central hypothesis is simple. If the reference dynamics generate unresolved information at rate h_{KS} while the tracking loop can constrain the reference only at effective rate C_{eff} , then the unresolved basis uncertainty should scale with the deficit

$$\kappa = h_{\text{KS}} - C_{\text{eff}} \ln 2.$$

In the chaos-wins regime ($\kappa > 0$), the predicted visibility loss has a characteristic double-exponential form and a breakdown time proportional to $1/\kappa$.

This is not presented as a completed theory of quantum measurement. It is a proposed operational bridge between finite-rate control theory and quantum visibility. A low-risk reference-channel demonstrator may deliberately throttle the basis-tracking channel and check whether online breakdown times collapse as a function of $h_{\text{KS}} - C_{\text{eff}} \ln 2$. That exercise is best understood as a calibration of the observer-relative control boundary. The higher-value physics test is the mesoscopic Penrose-overlap experiment, where the same variable can be compared against Penrose-style mass-geometry dependence on the same target platform.

The mesoscopic regime is strategically important because Penrose objective reduction (OR) predicts collapse-like timescales of the same order for certain masses and separations. The

two mechanisms are not the same: Penrose OR depends on gravitational self-energy, whereas the present proposal depends on reference-tracking bandwidth. Their overlap makes the regime experimentally valuable because a single setup may discriminate between mass-geometry scaling and tracking-capacity scaling.

Sections 2–3 define the reference-tracking variables and the finite-rate lemma; Sections 4–5 derive the visibility prediction; Section 6 presents testable predictions; Sections 7–8 compare with Penrose objective reduction and specify falsification criteria.

2 The Measurement Basis as Physical Variable

In standard treatments, the measurement basis θ is treated as a known external parameter. But in the laboratory θ is implemented by apparatus: a polarizer orientation, a magnetic field direction, a microwave pulse phase, a timing reference, or a local oscillator. The “choice” of θ is therefore also a physical process involving control electronics, calibration records, feedback loops, and finite precision.

This is a small reframing with a large consequence. The basis θ is not an external free parameter applied to the experiment from outside the world. It is a physical variable with its own causal history: produced by the controller, maintained by finite hardware, perturbed by chaotic or diffusive internal dynamics, and continually corrected by feedback. The setting and the system therefore share an ancestry in the same physical history. This does not require any new signal between them; it only acknowledges that the measurement question is itself an object of physics, not a free variable inserted from outside.

What follows from this is operational, not metaphysical. If the measurement basis is an internal physical variable, then maintaining it costs information processing rate, and losing track of it has experimental consequences. The standard treatment absorbs both costs into the idealization $\theta = \theta_0$. The present proposal asks under what conditions that idealization holds, and what visibility law replaces it when it does not. A philosophical and interpretive treatment of this move—and its connection to single-history ontologies, Bell-type constraints, and the broader IOF reading—is given in the accompanying document *Where Did the Measurement Basis Come From?*; the present paper restricts itself to the operational discriminator.

The proposal below does not require a new Bell model. It asks a more basic control question: under what conditions can the apparatus maintain and know its own reference frame well enough that the ideal basis parameter used in the quantum calculation remains experimentally valid? If finite-rate reference tracking contributes to visibility loss, the effect should be visible even in ordinary interferometric or qubit coherence experiments, without first resolving foundational debates about nonlocality.

The following sections develop a quantitative framework: we first define the reference-tracking capacity and the apparatus information-generation rate, then show how their interplay determines when tracking fails and what observable consequences follow.

3 The Tracking Problem

3.1 Definitions

We define an **observer/controller** as any physical system that:

1. Maintains or estimates a reference variable $\theta(t)$ used to define a measurement basis
2. Has finite effective channel capacity C_{eff} devoted to that reference-tracking task

The Landauer limit bounds the maximum physically possible information processing rate given available power P and temperature T :

$$C = \frac{P}{kT \ln 2} \quad (\text{bits/s}) \quad (1)$$

This is an upper bound, not an estimator of the tested quantity. In cryogenic electronics the Landauer bound can be enormous compared with the few-bit or few-hertz effective rate intentionally assigned to a particular tracking loop. The experimentally relevant C_{eff} is the information rate that actually constrains θ , preferably imposed directly by a digital rate limit or update-rate throttle. Power and temperature enter as thermodynamic consistency checks and as practical implementation knobs, not as substitutes for measuring or imposing C_{eff} .

3.2 Apparatus Dynamics: h_{KS}

We characterize the observer’s apparatus by its Kolmogorov-Sinai entropy rate h_{KS} —the information-production rate of the classical degrees of freedom (voltage references, timing circuits, feedback loops) that define and maintain the measurement basis.

Even deterministic feedback loops can exhibit chaotic sensitivity: tiny untracked differences grow exponentially, characterized by positive Lyapunov exponents and quantified by h_{KS} . Measurement apparatus naturally tends toward high h_{KS} because amplification from microscopic quantum events to macroscopic records requires nonlinear or threshold dynamics, which generically produce chaos. In contrast, purely diffusive dynamics ($h_{\text{KS}} \rightarrow 0$) produce only slow, linear growth of uncertainty—a limiting case where the framework predicts no capacity-dependent effects.

Quantum error correction can be understood as engineering that suppresses effective instability through finite-rate syndrome extraction, decoding, and feedback. The companion capacity-backaction analysis treats this as a standard quantum-engineering tradeoff: useful control capacity and measurement-induced physical instability must be accounted for together, without treating the quantum-computing case as an independent test of BLQC [8].

3.3 The Data-Rate Theorem: Scope and Application

The Data-Rate Theorem (DRT) establishes a necessary condition for mean-square stabilization: for a scalar unstable linear plant with Lyapunov exponent λ , stabilization over a finite-capacity channel requires $C > \lambda / \ln 2$ [1, 2]. For nonlinear systems, related results connect stabilizing data rates to topological and invariance entropy [2, 3].

The present application is narrower than a full nonlinear-control theorem. We use DRT logic to motivate a finite-rate phase-reference lemma that can be tested directly.

Toy lemma (finite-rate phase-reference tracking). Consider a one-dimensional reference error interval Δ_n for θ sampled every Δt . During one step, unresolved reference dynamics expand the interval by approximately $e^{\lambda \Delta t}$. A channel delivering $R \Delta t$ bits about the reference can reduce the interval by at most a factor $2^{-R \Delta t}$ under ideal coding. Thus, up to constants set by the coding scheme and noise model,

$$\Delta_{n+1} \sim e^{\lambda \Delta t} 2^{-R \Delta t} \Delta_n. \quad (2)$$

Taking the continuous-time limit gives

$$\frac{d}{dt} \ln \Delta \approx \lambda - R \ln 2. \quad (3)$$

Since variance scales as squared error amplitude,

$$\frac{d}{dt} \ln \sigma_\theta^2 \approx 2(\lambda - R \ln 2). \quad (4)$$

For a multi-mode chaotic reference, we replace λ by the relevant entropy-rate proxy h_{KS} and R by the effective tracking-channel rate C_{eff} . This replacement is the empirical bridge tested below. The proposal is not that Eq. (4) is a new theorem of quantum mechanics; it is that phase-reference visibility should reveal whether this finite-rate tracking law is operationally active for the online observer/controller, and whether that same coordinate remains predictive in the Penrose-overlap regime.

3.4 The Ignorance Rate

We define the **ignorance rate** κ as the gap between information generation and processing capacity:

$$\kappa = h_{\text{KS}} - C_{\text{eff}} \ln 2 \quad (5)$$

Following the DRT condition, this defines two regimes:

- **Capacity-wins** ($\kappa < 0$): Observer can in principle keep up with basis tracking
- **Chaos-wins** ($\kappa > 0$): Tracking error is predicted to grow until the basis is no longer resolved to the required tolerance

3.5 Operational Estimators

For the framework to be testable, C_{eff} and h_{KS} must be operationally defined.

Effective capacity C_{eff} : We define C_{eff} as the effective information rate in the basis-tracking loop (not the whole device):

$$C_{\text{eff}} = r \cdot b \cdot f \quad (6)$$

where r is the update rate (Hz), b is the effective bits per update constraining θ , and $f \in (0, 1]$ is the fraction of updates that actually constrain θ after overhead, latency, queuing, rejected samples, algorithmic processing limits, and filtering. In the cleanest experiment, C_{eff} is not inferred from total device power or raw Shannon bandwidth; it is imposed by a known rate limiter on the reference-tracking data stream or inferred from useful online estimator performance. In QEC-like settings the analogous capacity includes decoder throughput, clock-cycle budget, queue depth, and timely Pauli-frame update rate. Varying r or another useful-throughput parameter while holding temperature, readout power, plant dynamics, latency controls, and signal-to-noise ratio fixed is the preferred knob.

Metric entropy rate h_{KS} : For chaotic systems, h_{KS} equals the sum of positive Lyapunov exponents (Pesin identity). We estimate h_{KS} from the rate at which prediction error grows with prediction horizon: given controller state x_t , track how $|x_{t+n\Delta} - g^n(x_t)|$ increases with n , where g^n denotes n iterations of a fitted model. The exponential growth rate estimates λ_{max} , which provides a practical proxy (and lower bound) for h_{KS} . Standard algorithms apply to logged digital controller states (e.g., FPGA-based readout). The experiment should verify $h_{\text{KS}} > 0$ before testing capacity dependence.

4 From Phase-Reference Tracking to Visibility

The experimental claim does not require a hidden-variable model. It only requires a signal whose ideal quantum prediction depends on a physical reference phase or basis variable. Write the ideal interference contribution as

$$I_{\text{int}} \propto \cos(\phi - \theta), \quad (7)$$

where ϕ is the system phase and θ is the apparatus reference used to define the measurement basis.

The controller maintains an estimate $\hat{\theta}$ of the realized reference. The physical reference may differ from the intended value by

$$\theta = \theta_0 + \delta\theta, \quad (8)$$

where $\delta\theta$ is not assumed to be fundamental quantum noise. It is the unresolved tracking error of the reference loop. If the experimenter analyzes data as if $\theta = \theta_0$, this unresolved reference error appears as visibility loss.

The bridge hypothesis is therefore:

If the reference-tracking loop is in the chaos-wins regime, the variance of $\delta\theta$ follows the finite-rate law derived in Eq. (4), with λ replaced by the measured entropy-rate proxy h_{KS} and R by the imposed or estimated C_{eff} .

This framing makes the experiment conservative. A positive result would show that reference tracking contributes to visibility. A null result would rule out the proposed mechanism in the tested regime without requiring a verdict on broader interpretations of quantum mechanics.

5 Basis Uncertainty

5.1 Basis Tracking Error

Due to finite capacity, the controller's estimate $\hat{\theta}$ of the reference basis has uncertainty σ_{θ} . The realized $\theta(t_m)$ may differ from the intended value θ_0 . The key distinction is between the reference frame used in data analysis and the physical frame implemented by the apparatus. The experimenter records outcomes as if $\theta = \theta_0$, but the actual measurement may have been made in basis $\theta = \theta_0 + \delta\theta$.

Ordinary phase noise can also reduce visibility, so the mere presence of visibility loss would not support the model. The distinctive claim is the *scaling*: double-exponential decay with breakdown time $t_{\text{break}} \propto 1/\kappa$, where $\kappa = h_{\text{KS}} - C_{\text{eff}} \ln 2$ is built from independently measured or imposed control variables. If present, this structure could be missed when data are analyzed only with ordinary exponential or Gaussian dephasing models; this is a methodological concern rather than evidence that the effect has already been observed.

The framework does not require that capacity limits be uncircumventable in principle. It requires that *for a given online observer/controller with given capacity*, the limit exists and has predictable consequences. If a higher-capacity retrospective observer can recover the visibility from a high-resolution reference log, that recovery should be interpreted as observer-relative reference-frame loss, not irreversible physical decoherence. If C_{eff} increases, t_{break} should increase accordingly (Eq. 14). The claim is that *wherever* the capacity limit lies, it determines

online visibility loss via κ . Whether one labels this “fundamental” or “technical” is mostly a terminological choice for a standalone demonstrator; the Penrose-overlap regime is where the distinction becomes physically decisive.

5.2 Dynamics

In the chaos-wins regime ($\kappa > 0$), we model the basis estimation error as inheriting the exponential separation rate of the uncontrolled reference dynamics, reduced by the effective information rate of corrective updates. For a chaotic mode with Lyapunov exponent λ , uncorrected estimation error scales as $e^{\lambda t}$; therefore its *variance* scales as $e^{2\lambda t}$. Finite-rate correction reduces the effective exponent by $C_{\text{eff}} \ln 2$, yielding:

$$\frac{d}{dt} \ln \sigma_\theta^2 \approx 2(h_{\text{KS}} - C_{\text{eff}} \ln 2) = 2\kappa \quad (9)$$

This is the finite-rate tracking hypothesis on which the paper rests. It is the minimal application of data-rate logic to observer-side basis tracking: instability expands unresolved basis uncertainty, finite capacity contracts it, and the net observable rate is their difference. The proposed experiment tests exactly this scaling.

Integrating Eq. (9):

$$\sigma_\theta^2(t) = \sigma_0^2 e^{2\kappa t} \quad (10)$$

The **threshold-crossing time** for uncertainty to grow from σ_0^2 to tolerance σ_{tol}^2 is:

$$\tau_{\text{loss}} = \frac{1}{2\kappa} \ln \left(\frac{\sigma_{\text{tol}}^2}{\sigma_0^2} \right) \quad (\kappa > 0) \quad (11)$$

5.3 Visibility Decay

When the observer has Gaussian uncertainty in its measurement basis, the observed visibility (fringe contrast) decays as:

$$V_{\text{IOF}}(t) = \exp \left(-\frac{\sigma_0^2}{2} e^{2\kappa t} \right) \quad (12)$$

where σ_0^2 is the initial basis uncertainty at $t = 0$, set by apparatus calibration and initialization quality.

Validity: This assumes Gaussian angle uncertainty and is accurate for $\sigma_\theta \lesssim 1$ rad. For larger uncertainties, higher cumulants contribute and suppression saturates.

Derivation: The observer intends basis θ_0 but realizes $\theta = \theta_0 + \delta\theta$ where $\delta\theta \sim \mathcal{N}(0, \sigma_\theta^2)$. Averaging the interference term $\cos(\phi - \theta)$ over this distribution:

$$\langle \cos(\phi - \theta) \rangle = \cos(\phi - \theta_0) \cdot e^{-\sigma_\theta^2/2} \quad (13)$$

5.4 Observable Breakdown Time

Define $V_* \in (0, 1)$ as a visibility threshold. The **breakdown time** t_{break} is when $V(t_{\text{break}}) = V_*$:

$$t_{\text{break}} = \frac{1}{2\kappa} \ln \left(\frac{-2 \ln V_*}{\sigma_0^2} \right) \quad (\kappa > 0) \quad (14)$$

This is the primary experimental observable.

6 Correlation Attenuation and No-Signaling Check

The core experiment is an interferometric visibility test, not a Bell-theorem replacement. If the same reference-tracking mechanism is tested in a Bell-type platform, its immediate prediction is conservative: basis uncertainty attenuates correlations below the ideal quantum value while preserving single-party marginals.

With Gaussian basis uncertainty σ^2 that is unbiased and outcome-independent, marginals remain exactly 50/50 (no signaling), but joint correlations are attenuated:

$$\langle E \rangle = -\cos(\theta_A - \theta_B) \cdot e^{-(\sigma_A^2 + \sigma_B^2)/2} \quad (15)$$

This is the independent-error form (independent Gaussian basis errors on the two wings); with correlated reference drift the exponent becomes $-\frac{1}{2} \text{Var}(\delta\theta_A - \delta\theta_B)$, recovering the above when the cross-covariance vanishes. As uncertainty grows, the CHSH parameter $|S|$ decreases from $2\sqrt{2}$ toward zero—correlations fall *below* ideal QM predictions. Any systematic change in one-party marginals should be treated as an instrumental bias or failed control condition, not as support for the present model.

The novel prediction is the κ -dependent scaling: $\sigma^2(t) = \sigma_0^2 e^{2\kappa t}$, yielding double-exponential visibility decay $V_{\text{IOF}}(t) = \exp(-\frac{1}{2}\sigma_0^2 e^{2\kappa t})$ with breakdown time $t_{\text{break}} \propto 1/\kappa$. Here V_{IOF} is the capacity-dependent IOF factor only; the observed visibility is the two-channel product $V_{\text{obs}} = V_{\text{std}} V_{\text{IOF}}$, with V_{std} the ordinary decoherence the framework does not modify. Figure 1 illustrates this.

7 Testable Predictions

The framework makes specific predictions that can be compared with standard dephasing and standard online-control failure. The cleanest reference-channel demonstrator is deliberate throttling of the reference-tracking channel. At fixed temperature, fixed readout signal-to-noise ratio, fixed plant dynamics, fixed sequence duration, and fixed environmental coupling, increasing the imposed tracking rate C_{eff} should delay online breakdown:

$$\frac{\partial t_{\text{break}}}{\partial C_{\text{eff}}} > 0 \quad (\kappa > 0). \quad (16)$$

More specifically, the measured breakdown times should collapse against

$$\kappa = h_{\text{KS}} - C_{\text{eff}} \ln 2, \quad (17)$$

with

$$t_{\text{break}} = \frac{1}{2\kappa} \ln \left(\frac{-2 \ln V_*}{\sigma_0^2} \right). \quad (18)$$

Power is a secondary implementation knob. If increasing controller power increases the effective tracking rate while temperature and readout SNR are clamped, the model predicts a sign reversal relative to ordinary heating: extra useful online capacity should extend observer-relative visibility time rather than shorten it. If power is increased without increasing the imposed or measured C_{eff} , no improvement is predicted. This distinction is important because the Landauer

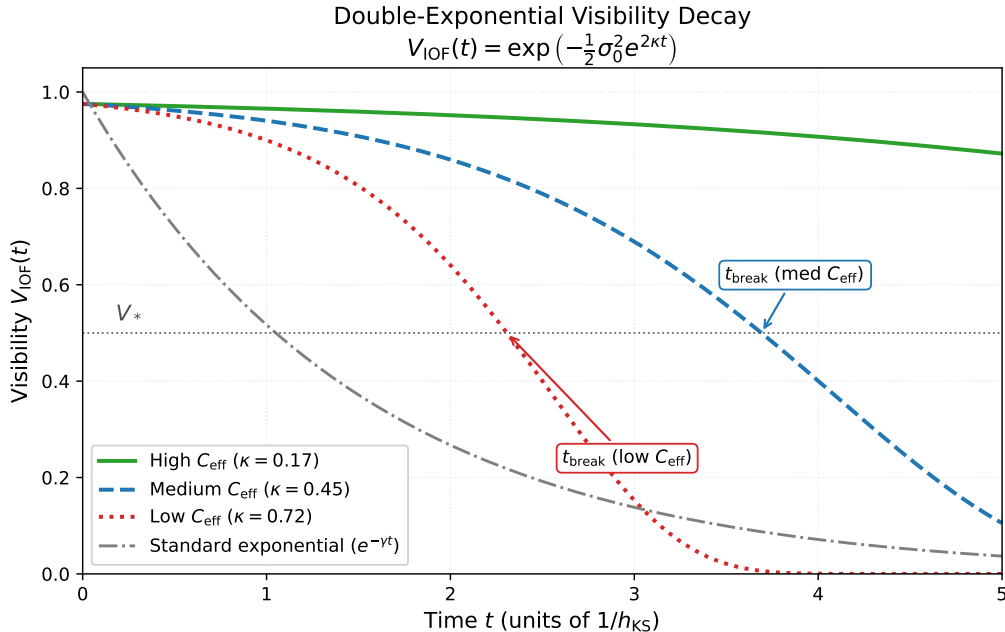


Figure 1: **Double-exponential visibility decay for three capacity levels.** Visibility $V_{\text{IOF}}(t) = \exp(-\frac{1}{2}\sigma_0^2 e^{2\kappa t})$ is plotted for high capacity (green, $\kappa = 0.17$), medium capacity (blue, $\kappa = 0.45$), and low capacity (red, $\kappa = 0.72$). Lower observer capacity C_{eff} yields higher $\kappa = h_{\text{KS}} - C_{\text{eff}} \ln 2$, causing earlier breakdown. The arrows mark t_{break} , the time when visibility crosses the threshold $V_* = 0.5$. For comparison, standard exponential decoherence ($e^{-\gamma t}$, gray dash-dot) decays steadily from $t = 0$, whereas the double-exponential curves remain near unity before falling sharply—a qualitatively distinct signature. The prediction: reducing C_{eff} shifts t_{break} leftward.

expression is only an upper bound on possible processing, not the actual rate assigned to the basis-reference loop.

Observer-relative visibility should follow the double-exponential decay $V_{\text{IOF}}(t) = \exp(-\frac{1}{2}\sigma_0^2 e^{2\kappa t})$, qualitatively distinct from standard exponential ($e^{-\gamma t}$) or Gaussian (e^{-t^2}) decoherence. Higher apparatus entropy rate h_{KS} should produce faster online visibility loss (smaller t_{break}), independent of environmental decoherence channels, when C_{eff} is held fixed. A fixed-duration loss that is fully recoverable from a withheld high-resolution reference log is a successful observer-relative tracking benchmark, but not evidence of irreversible physical collapse.

Qualitative distinctness is not the same as measurability. With finite shot noise and a finite observation window, the double-exponential bend must be separated not from a fixed exponential but from the *closest* ordinary-decay curve. Figure 2 quantifies this for the strongest such single adversary, the best-fitting stretched exponential $\exp[-(\Gamma t)^\beta]$, which contains pure-exponential ($\beta = 1$) and Gaussian ($\beta = 2$) decay as special cases. Modelling each visibility estimate with fringe shot noise (variance $\approx (1 - V^2)/N$ per point) and refitting both forms to every dataset, the receiver-operating-characteristic area (AUC) for distinguishing them rises from chance (0.5) to near-certain (1) only once the observation window passes t_{break} : below roughly half of t_{break} the bend has not yet developed and the two forms are statistically indistinguishable, while past the knee a few hundred shots per point at percent-level fringe precision suffice for clean separation. The operational consequence is that the wrong-functional-form falsification criterion stated below has discriminating power only when the window reaches the breakdown knee; an experiment that stops short of t_{break} cannot separate the IOF signature from ordinary dephasing at any shot count. This is a statement about measurement budget, not evidence for the framework; the full method and code accompany the online verification materials.

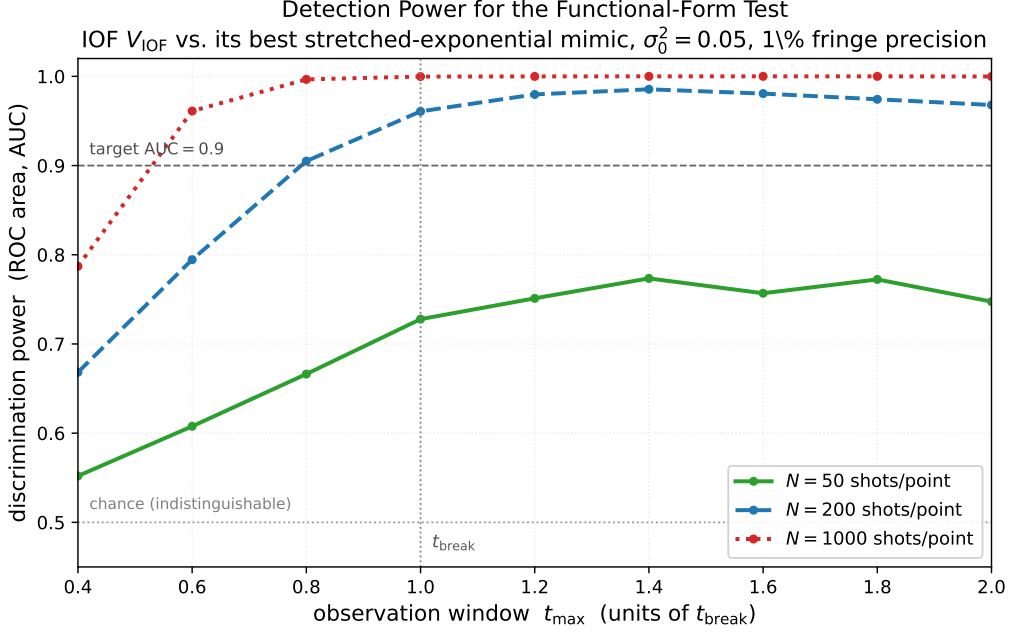


Figure 2: **Detection power for the functional-form test.** ROC area (AUC) for distinguishing the IOF double-exponential $V_{\text{IOF}}(t) = \exp(-\frac{1}{2}\sigma_0^2 e^{2\kappa t})$ from its best stretched-exponential mimic $\exp[-(\Gamma t)^\beta]$, versus the observation window t_{\max} (in units of t_{break}), for three shot budgets per time point (green $N = 50$, blue $N = 200$, red $N = 1000$; $\sigma_0^2 = 0.05$, 1% fringe precision). AUC = 0.5 is chance (statistically indistinguishable); AUC \rightarrow 1 is clean separation. The signature becomes resolvable only as the window passes t_{break} , and beyond a few hundred shots measurement precision—not shot count—sets the limit. The axis is normalised to t_{break} , so the curve is independent of κ and of the absolute timescale. A statement about experimental budget, not evidence for the framework.

A comprehensive experimental protocol specifying platform requirements, statistical analysis plans, and detailed procedures is available separately [4]. QGEM-style pathfinder apparatuses [7], designed for extreme isolation and cryogenic operation with mesoscopic masses, may provide useful future platforms if they expose the required reference-tracking and mass-geometry knobs.

8 Penrose-Overlap Regime

The framework’s predictions can be compared with other approaches to quantum state reduction. Penrose’s gravitational Objective Reduction (OR) provides a natural comparison because both predict characteristic waveform-collapse or visibility-loss timescales in a mesoscopic regime, while depending on different physical variables. This overlap is not used as evidence for the present mechanism. It is used to identify a high-value experimental regime in which two candidate mechanisms can be discriminated.

Penrose proposes that quantum superpositions become unstable due to gravitational self-energy, giving a characteristic timescale [5]:

$$\tau_{\text{OR}} = \frac{\hbar}{E_G} \sim \frac{\hbar s}{Gm^2} \quad (19)$$

where m is the mass in superposition and s is the spatial separation. For mesoscopic masses ($m \sim 10^{-15}$ kg) with separations of 100 nm–1 μm , this gives $\tau_{\text{OR}} \sim 10$ –100 ms.

The present framework predicts a tracking-loss timescale in the chaos-wins regime:

$$\tau_{\text{loss}} \approx \frac{1}{\kappa} = \frac{1}{h_{\text{KS}} - C_{\text{eff}} \ln 2} \quad (20)$$

For apparatus parameters $h_{\text{KS}} \approx 50$ nats/s and $C_{\text{eff}} \approx 10$ bits/s ($C_{\text{eff}} \ln 2 \approx 7$ nats/s), we obtain $\kappa \approx 43$ nats/s and $\tau_{\text{loss}} \approx 1/\kappa \approx 23$ ms. Including the log factor from Eq. (11) for 1–5% visibility thresholds contributes a factor of 2–3, giving $\tau_{\text{loss}} \approx 50$ –70 ms.

Numerical proximity: For the apparatus parameters above, τ_{loss} falls near the Penrose-predicted τ_{OR} range for mesoscopic systems. This is the “Penrose-overlap” regime: both mechanisms can predict loss on comparable timescales, but they assign the timescale to different knobs.

Different predictions: Despite this proximity, the frameworks make predictions with different primary dependencies:

- Penrose OR: gravitational instability; τ_{OR} depends on mass and geometry, independent of observer bandwidth
- Present framework: information-theoretic tracking limits; τ_{loss} depends on useful online reference-tracking capacity, independent of spatial geometry

This enables experimental discrimination:

Test	Penrose OR	This framework
Vary C_{eff} at fixed mass/geometry	No effect on τ	τ_{loss} increases
Vary h_{KS} at fixed mass/geometry	No effect on τ	τ_{loss} decreases
Vary separation s at fixed C_{eff}	τ_{OR} increases	No effect on τ_{loss}

The primary discriminator is therefore not the absolute collapse time, but the derivative of that time with respect to independently controlled variables. If increasing useful reference-tracking capacity extends the relevant visibility or coherence time while mass geometry is fixed, the information-theoretic mechanism is supported in that regime. If mass-geometry variables explain the data and κ adds no predictive value after ordinary confounds are controlled, gravitational OR remains the better explanation in this regime. The two mechanisms are not framed as mutually exclusive: the full statistical treatment of the overlap regime—including the additive combined-rate model, geometry–capacity mediation, and the conditional bridge to the speculative *Bridge Ansatz* of [9]—is developed in Appendix B of the experimental protocol [4].

Penrose OR is considered foundational because it connects state reduction to gravity—a universal physical phenomenon. The present framework instead connects visibility loss to information-theoretic limits on physical observers, which are also universal in scope because all physical control systems are subject to thermodynamic constraints such as the Landauer bound. Thermodynamic approaches to gravity—notably Jacobson’s derivation of Einstein’s equations from local horizon thermodynamics [6]—suggest these domains may be connected more deeply than their separate formulations imply. If gravitational and information-theoretic scales are linked, the numerical proximity of τ_{loss} and τ_{OR} might reflect structure rather than coincidence; at present this remains a speculative motivation for discriminating experiments, not an established conclusion.

9 Falsification Criteria

The framework is falsified if any of the following hold:

1. **Wrong capacity dependence:** At constant temperature, fixed readout SNR, fixed sequence duration, and independently monitored environmental channels, increasing the imposed reference-tracking rate C_{eff} fails to extend online breakdown time, i.e., $\partial t_{\text{break}}/\partial C_{\text{eff}} \leq 0$ in the claimed chaos-wins regime.
2. **No data collapse:** Breakdown times fail to collapse against $h_{\text{KS}} - C_{\text{eff}} \ln 2$ after h_{KS} and C_{eff} are independently measured or imposed.
3. **No chaos dependence:** t_{break} is independent of apparatus chaos rate h_{KS}
4. **Wrong functional form:** $V_{\text{IOF}}(t)$ fits ordinary exponential or power-law decoherence but *not* the double-exponential structure $V_{\text{IOF}}(t) = \exp(-\frac{1}{2}\sigma_0^2 e^{2\kappa t})$
5. **Gravitational scaling in the Penrose-overlap regime:** Coherence timescale tracks mass/geometry ($\propto m^{-2}$ or $\propto s$) and κ adds no predictive value after ordinary confounds are controlled—gravitational OR explains the data in this regime
6. **Marginal violation:** Single-party statistics deviate from 50/50, indicating signaling or systematic bias

10 Conclusion

We have presented a finite-rate phase-reference test. The measurement basis is treated as a physical reference variable $\theta(t)$ that must be maintained by a finite-rate tracking loop. When the reference entropy rate exceeds effective tracking capacity, the model predicts exponential growth of basis uncertainty.

The key prediction is operational: in the chaos-wins regime, online observer-relative visibility decays as $V_{\text{IOF}}(t) = \exp(-\frac{1}{2}\sigma_0^2 e^{2\kappa t})$, with breakdown time inversely proportional to $\kappa = h_{\text{KS}} - C_{\text{eff}} \ln 2$. Controlled variation of C_{eff} and h_{KS} provides a direct benchmark of the finite-rate control boundary. These predictions are falsifiable without assuming a hidden-variable model or a new interpretation of Bell correlations.

The framework’s timescale $\tau_{\text{loss}} \approx 1/\kappa$ can overlap with Penrose’s τ_{OR} in the mesoscopic regime ($\sim 50\text{--}70$ ms), but predicts different primary dependencies—tracking bandwidth and reference entropy rate versus mass geometry. That overlap makes the regime high-value: a single experiment can constrain the joint contributions of finite-rate reference tracking and gravitational objective reduction by varying the appropriate knobs. The two mechanisms are not assumed mutually exclusive; Appendix B of [4] treats additive, mediated, and collinear (Bridge-Ansatz) outcomes within the same regression.

If the predicted capacity dependence is observed in a reference-channel demonstrator, basis-reference tracking becomes a calibrated component of observer-relative quantum visibility. If the same coordinate remains predictive in the mesoscopic Penrose-overlap regime after ordinary latency, pulse-error, wait-time, thermal, and offline-recoverable tracking artifacts are controlled, the proposed bridge becomes physically nontrivial. If it is not observed, the proposed bridge fails there.

References

- [1] S. Tatikonda and S. Mitter, “Control under communication constraints,” *IEEE Trans. Autom. Control* **49**, 1056–1068 (2004).
- [2] G.N. Nair and R.J. Evans, “Stabilizability of stochastic linear systems with finite feedback data rates,” *SIAM J. Control Optim.* **43**, 413–436 (2004).
- [3] C. Kawan, *Invariance Entropy for Deterministic Control Systems*, Lecture Notes in Mathematics 2089, Springer (2013).
- [4] A. Dekker, “Prospective Tests of Bandwidth-Limited Observer Dynamics: Experimental Protocol for the Ignorant Observer Framework,” OSF Preprints, v2.5 (2026). DOI: 10.17605/OSF.IO/2QJNE
- [5] R. Penrose, “On Gravity’s Role in Quantum State Reduction,” *Gen. Relativ. Gravit.* **28**, 581–600 (1996).
- [6] T. Jacobson, “Thermodynamics of Spacetime: The Einstein Equation of State,” *Phys. Rev. Lett.* **75**, 1260–1263 (1995).
- [7] S. Bose, A. Mazumdar, M. Schut, and M. Toroš, “Mechanism for the quantum natured gravitons to entangle masses,” *Phys. Rev. D* **105**, 106028 (2022); QGEM Collaboration, “Pathfinder design for tabletop tests of quantum gravity,” arXiv:2509.01586 (2025).
- [8] A. Dekker, “The Capacity–Backaction Frontier: Aggregate Rate Accounting for Fault-Tolerant Quantum Control,” manuscript (2026).
- [9] A. Dekker, “The Creation of Duality: A Speculative Extension on Appearance, Gravity, and Information,” OSF Preprints, v1.3 (2026). DOI: 10.17605/OSF.IO/C5EN8.